REMARKS

Claims 1-7, 10-13, 16-20, 22, 24, and 26 are pending in the application.

Claims 1-7, 10-13, 16-20, 22, 24, and 26 are currently amended, and claims 8, 9, 14, 15, 21, 23, 25, and 27 are cancelled Applicants respectfully submit that no new matter is added to currently amended claims 1-7, 10-13, 16-20, 22, 24, and 26.

Applicant respectfully submits that entry of the currently amended claims is proper because the currently amended claims will either place the application in condition for allowance or in better form for appeal.

Claims 1-8, 10-13, 16-18 and 23 stand rejected under 35 U.S.C. §102(b) as anticipated by "The Mirage NFS Router" by Baker et al., a technical report published by University of Arizona, 2002, hereinafter, Baker.

Claim 9 and 19-22 stand rejected under 35 U.S.C. §103(a) as unpatentable over Baker in view of "NAS Switch: A Novel CIFS Server Virtualization" by Katsurashima et al., 2003, hereinafter, Katsurashima.

Claims 1-13, 16-22, and 24-26 stand rejected under 35 U.S.C. §103(a) as unpatenable over U.S. Patent Application Publication No. 2002/0120763 to Miloushev et al., hereinafter, Miloushev in view of "Network File System Protocol Specification", IETF RFC 1094, version 2.0, hereinafter, RFC 1094.

Claims 1-13, 16-22, and 24-26 stand rejected under 35 U.S.C. §112, first paragraph.

Applicant respectfully traverses the rejections based on the following discussion.

I.   The Prior Art Rejections
A.   The 35 U.S.C. 102(b) Rejection over Baker
1.   The Baker Disclosure

Baker discloses a system (Mirage) that aggregates multiple NFS servers into a single, virtual NFS file server. It is interposed between the NFS clients and servers, making the clients believe that they are communication with a single large server. Mirage is an NFS router because it routes an NFS request from a client to the proper NSF server, and routes the reply back to the

proper client. Mirage also prevents DoS atttacks on the NFS protocol, ensuring that all clients receive a fair share of the servers' resources. Mirage is designed to run on an IP router, providing virtualized NFS file service without affecting other network traffic. Experiments with a Mirage prototype show that Mirage effectively virtualizes an NFS server using unmodified clients and servers, and it ensures that legitimate clients receive a fair share of the NFS server even during a DoS attack. Mirage imposes an overhead of only 7% on a realistic NFS workload. (Abstract)

2. Arguments

Currently amended, independent claims 1 and 12 recite in relevant part,

"a communication virtualizer,

wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet,

wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and

wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol".

Baker merely discloses a system that aggregates multiple NFS servers into a single, virtual NFS file server.

Nowhere does Baker disclose, teach or suggest combining multiple standard Ethernet packets, whose data size exceeds that of a maximum size supported by a NFS protocol, into a jumbo packet.

Instead, Baker merely discloses a system that aggregates multiple NFS servers into a single, virtual NFS file server.

For at least the reasons outlined above, Applicants respectfully submit that Baker does not disclose, teach or suggest at least the present invention's features of: "a communication virtualizer, wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet, wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File

System (NFS) protocol, and wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol", as recited in currently amended, independent claims 1 and 12. Accordingly, Baker does not anticipate the subject matter of currently amended, independent claims 1 and 12, and dependent claims 2-7, 10, 11, 13, and 16-18 under 35 U.S.C. §102(b). The rejection of cancelled claims 8 and 23 is moot. Withdrawal of the rejection of claims 1-8, 10-13, 16-18 and 23 under 35 U.S.C. §102(b) as anticipated by Baker is respectfully solicited.

B. The 35 U.S.C. 103(a) Rejection over Baker and Katsurashima

1. The Katsurashima Disclosure

Katsurashima discloses a Common Internet File System (CIFS) server virtualization method which requires no proprietary software or hardware for clients or NAS units. The method is implemented as an in-band network application between clients and NAS units, and it provides users and administrators with a single virtual NAS system that incorporates all their units. Since almost all name resolution operations are performed by individual NAS units independently, use of this method imposes only a very light computational load and creates little latency. (Abstract).

2. Arguments

Katsurashima does not cure the deficiencies of Baker.

Katsurashima merely discloses a Common Internet File System (CIFS) server virtualization method.

Nowhere does Katsurashima disclose, teach or suggest combining multiple standard Ethernet packets, whose data size exceeds that of a maximum size supported by a NFS protocol, into a jumbo packet.

Instead, Katsurashima merely discloses a Common Internet File System (CIFS) server virtualization method.

Applicants respectfully submit that nowhere does Katsurashima disclose, teach or suggest at least the present invention's features of: "a communication virtualizer, wherein said

communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet, wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol ", as recited in currently amended, independent claim 12.

For at least the reasons outlined above with respect to the rejection of the claims over Baker and for at least those reasons outlined immediately above with respect to the rejection of the claims over Katsurashima, Applicants respectfully submit that Baker and Katsurashima, either individually or in combination, do not disclose, teach or suggest at least the present invention's features of: "a communication virtualizer, wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet, wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol ", as recited in currently amended, independent claim 12. Accordingly, Baker and Katsurashima, either individually or in combination, fail to render obvious the subject matter of currently amended, independent claim 12, and dependent claims 19, 20, and 22 under 35 U.S.C. § 103(a). The rejection of cancelled claims 9 and 21 is moot. Withdrawal of the rejection of claims 9 and 19-22 under 35 U.S.C. § 103(a) as unpatentable over Baker and Katsurashima is respectfully solicited.

C.   The 35 U.S.C. 103(a) Rejection over Miloushev and RFC 1094

1.   The Miloushev Disclosure

Miloushev discloses that an apparatus and method are provided in a computer network to decouple the client from the server, by placing a transparent network node, also termed a file switch or file switch computer, between the client and the server. Usage of such a file switch allows reduced latency in file transfers, as well as scalable mirroring, striping, spillover, and other features. (Abstract).

2. The RFC 1094 Disclosure

RFC 1094 discloses a Sun Network Filesystems (NFS) protocol that provides transparent remote access to shared files across networks. The NSF protocol is designed to be portable across different machines, operating systems, network architectures, and transport protocols. This portability is achieved through the use of Remote Procedure Call (RPC) primitives built on top of an eXternal Data Representation (XDR). Implementations already exist for a variety of machines, from personal computers to supercomputers.

3. Arguments

Currently amended, independent claims 1 and 12 recite in relevant part,

"a communication virtualizer,

wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet,

wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and

wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol ".

Miloushev merely discloses that an apparatus and method are provided in a computer network to decouple the client from the server, by placing a transparent network node, also termed a file switch or file switch computer, between the client and the server.

Nowhere does Miloushev disclose, teach or suggest combining multiple standard Ethernet packets, whose data size exceeds that of a maximum size supported by a NFS protocol, into a jumbo packet.

Instead, Miloushev merely discloses that an apparatus and method are provided in a computer network to decouple the client from the server, by placing a transparent network node, also termed a file switch or file switch computer, between the client and the server.

For at least the reasons outlined above, Applicants respectfully submit that Miloushev does not disclose, teach or suggest at least the present invention's features of: "a communication

virtualizer, wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet, wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol ", as recited in currently amended, independent claims 1 and 12.

RFC 1094 does not cure the deficiencies of Miloushev.

Nowhere does RFC 1094 disclose, teach or suggest combining multiple standard Ethernet packets, whose data size exceeds that of a maximum size supported by a NFS protocol, into a jumbo packet.

Instead, RFC 1094 merely discloses a Sun Network Filesystems (NFS) protocol that provides transparent remote access to shared files across networks.

For at least the reasons outlined above with respect to the rejection of the claims over Miloushev and for at least those reasons outlined immediately above with respect to the rejection of the claims over RFC 1094, Applicants respectfully submit that Miloushev and RFC 1094, either individually or in combination, do not disclose, teach or suggest at least the present invention's features of: "a communication virtualizer, wherein said communication virtualizer combines multiple standard Ethernet packets received from a client computer into a jumbo packet, wherein a data size of said multiple standard Ethernet packets received exceeds that of a maximum size supported by a Network File System (NFS) protocol, and wherein a data size of said jumbo packet exceeds that of said maximum size supported by said NFS protocol", as recited in currently amended, independent claims 1 and 12. Accordingly, Miloushev and RFC 1094, either individually or in combination, fail to render obvious the subject matter of currently amended, independent claims 1 and 12, and dependent claims 2-7, 10-13, 16-20, 22, 24, and 16 under 35 U.S.C. §103(a). The rejection of cancelled claims 8, 9, 21, and 25 is moot. Withdrawal of the rejection of claims 1-13, 16-22, and 24-26 under 35 U.S.C. §103(a) as unpatentable over Miloushev and RFC 1094 is respectfully solicited.

II.   The 35 U.S.C. §112, First Paragraph, Rejection

Claims 1-13, 16-22, and 24-26 stand rejected under 35 U.S.C. §112, first paragraph, because the Final Action asserts that an element, "a client indentifer", was not included in the originally filed application.

Applicants respectfully submit that the phrase, "a client identifier", has been deleted from the currently amended claims, above.

Applicants have also amended claim 12 to recite in relevant part, "wherein each of said response packets identifies said client computer", which receives support from paragraph [0038], lines 1-3 of the originally filed Specification.

For at least the reasons outlined above, Applicants respectfully submit that the currently amended claims fulfill the statutory requirements of 35 U.S.C. §112, first paragraph. Withdrawal of the rejection of claims 1-13, 16-22, and 24-26 under 35 U.S.C. §112, first paragraph, is respectfully solicited.

III.   Formal Matters and Conclusion

Claims 1-7, 10-13, 16-20, 22, 24, and 26 are pending in the application.

Applicant respectfully submits that entry of currently amended claims 1-7, 10-13, 16-20, 22, 24, and 26 is proper because the currently amended claims will either place the application in condition for allowance or in better form for appeal.

With respect to the rejections of the claims over the cited prior art, Applicants respectfully argue that the currently amended claims are distinguishable over the prior art of record. In view of the foregoing, the Examiner is respectfully requested to reconsider and withdraw the prior art rejections to the claims.

In view of the foregoing, Applicants submit that claims 1-7, 10-13, 16-20, 22, 24, and 26, all the claims presently pending in the application, are in condition for allowance. The Examiner is respectfully requested to pass the above application to issue at the earliest time possible.

Should the Examiner find the application to be other than in condition for allowance, the Examiner is requested to contact the undersigned at the local telephone number listed below to discuss any other changes deemed necessary.

Please charge any deficiencies and credit any overpayments to Attorney's Deposit Account Number 09-0441.

Respectfully submitted,

Dated: May 12, 2008

/Peter A. Balnave/
Peter A. Balnave, Ph.D.
Registration No. 46,199

Gibb & Rahman, LLC
2568-A Riva Road, Suite 304
Annapolis, MD 21401
Voice: (410) 573-5255
Fax: (301) 261-8825
Email: Balnave@Gibb-Rahman.com
Customer Number: 29154